ECEN 615
Methods of Electric Power Systems Analysis

Lecture 8: Sparse Systems

Prof. Tom Overbye
Dept. of Electrical and Computer Engineering
Texas A&M University
overbye@tamu.edu

Announcements

- Read Chapter 6 from the book
  - The book presents the power flow using the polar form for the Y_{bus} elements
- Homework 2 is due on Thursday September 22
- The ECEN EPG dinner is Saturday October 1st at 5:00 pm at Prof. Davis's house. The address is 3810 Park Meadow Lane, Bryan, TX 77802.
  - **RSVP** at **https://forms.gle/6aye2butgCLDv6bz5** **by Sunday, September 25th.**

LU Algorithm Without Pivoting Processing by row

- We will use the more common approach of having ones on the diagonals of **L**. Also in the common, diagonally dominant power system problems pivoting is not needed.
  The below algorithm is in row form (useful with sparsity!)

```
For i := 2 to n Do Begin  // This is the row being processed
  For j := 1 to i-1 Do Begin  // Rows subtracted from row i
    A[i,j] = A[i,j]/A[j,j]  // This is the scaling
    For k := j+1 to n Do Begin  // Go through each column in i
      A[i,k] = A[i,k] - A[i,j]*A[j,k]
    End;
  End;
End;
```

Forward Substitution

Forward substitution solves **b = Ly** with values in **b** being over written (replaced by the **y** values)

```
For i := 2 to n Do Begin  // This is the row being processed
  For j := 1 to i-1 Do Begin
    b[i] = b[i] - A[i,j]*b[j]    // This is just using the L matrix
  End;
End;
```

Backward Substitution

- Backward substitution solves $\mathbf{y} = \mathbf{U}\mathbf{x}$ (with values of $\mathbf{y}$ contained in the $\mathbf{b}$ vector as a result of the forward substitution)

```
For i := n to 1 Do Begin  // This is the row being processed
  For j := i+1 to n Do Begin
    b[i] = b[i] - A[i,j]*b[j]    // This is just using the U matrix
  End;
  b[i] = b[i]/A[i,i]    // The A[i,i] values are <> 0 if it is nonsingular
End
```

Computational Complexity

- Computational complexity indicates how the number of numerical operations scales with the size of the problem
- Computational complexity is expressed using the "Big O" notation; assume a problem of size n
  - Adding the number of elements in a vector is O(n)
  - Adding two n by n full matrices is $O(n^2)$
  - Multiplying two n by n full matrices is $O(n^3)$
  - Inverting an n by n full matrix, or doing Gaussian elimination is $O(n^3)$
  - Solving the traveling salesman problem by brute-force search is O(n!)

Computational Complexity

- Knowing the computational complexity of a problem can help to determine whether it can be solved (at least using a particular method)
  - Scaling factors do not affect the computation complexity
    - an algorithm that takes $n^3/2$ operations has the same computational complexity of one the takes $n^3/10$ operations (though obviously the second one is faster!)
- With $O(n^3)$ factoring a full matrix becomes computationally intractable quickly!
  - A 100 by 100 matrix takes a million operations (give or take)
  - A 1000 by 1000 matrix takes a billion operations
  - A 10,000 by 10,000 matrix takes a trillion operations!

Sparse Systems

- The material presented so far applies to any arbitrary linear system
- The next step is to see what happens when we apply triangular factorization to a sparse matrix
- For a sparse system, only nonzero elements need to be stored in the computer since no arithmetic operations are performed on the 0's
- The LU factorization is adapted to solve sparse systems in such a way as to preserve the sparsity as much as possible

Sparse Matrix History

- A nice overview of sparse matrix history is by Iain Duff at http://www.siam.org/meetings/la09/talks/duff.pdf
- Sparse matrices developed simultaneously in several different disciplines in the early 1960's with power systems definitely one of the key players (Bill Tinney from BPA)
- Different disciplines claim credit since they didn't necessarily know what was going on in the others

Sparse Matrix History

- In power systems a key N. Sato, W.F. Tinney, “Techniques for Exploiting the Sparsity of the Network Admittance Matrix,” Power App. and Syst., pp 944-950, December 1963
  - In the paper they are proposing solving systems with up to 1000 buses (nodes) in 32K of memory!
  - You’ll also note that in the discussion by El-Abiad, Watson, and Stagg they mention the creation of standard test systems with between 30 and 229 buses (this surely included the now famous 118 bus system)
  - The BPA authors talk “power flow” and the discussors talk “load flow.”
- Tinney and Walker present a much more detailed approach in their 1967 IEEE Proceedings paper titled “Direct Solutions of Sparse Network Equations by Optimally Order Triangular Factorization”

Sparse Matrix Computational Order

- The computational order of factoring a sparse matrix, or doing a forward/backward substitution depends on the matrix structure
  - Full matrix is $O(n^3)$
  - A diagonal matrix is O(n); that is, just invert each element
- For power system problems the classic paper is F. L. Alvarado, "Computational complexity in power systems," *IEEE Transactions on Power Apparatus and Systems*, ,May/June 1976
  - $O(n^{1.4})$ for factoring, $O(n^{1.2})$ for forward/backward
  - For a 100,000 by 100,000 matrix changes computation for factoring from 1 quadrillion to 10 million!
  - Updated statistics are given in F. Safdarian, Z. Mao, W. Jang, and T. J. Overbye, "Power System Sparse Matrix Statistics," IEEE Texas Power and Energy Conference, College Station, TX, February 2022

Inverse of a Sparse Matrix

- The inverse of a sparse matrix is NOT in general a sparse matrix
- We never (or at least very, very, very seldom) explicitly invert a sparse matrix
  - Individual columns of the inverse of a sparse matrix can be obtained by solving $\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$ with **b** set to all zeros except for a single nonzero in the position of the desired column
  - If a few desired elements of $\mathbf{A}^{-1}$ are desired (such as the diagonal values) they can usually be computed quite efficiently using sparse vector methods (a topic we'll be considering soon)
- We can't invert a singular matrix (whether sparse or full)

Computer Architecture Impacts

- With modern computers the processor speed is many times faster than the time it takes to access data in main memory
  - Some instructions can be processed in parallel
- Caches are used to provide quicker access to more commonly used data
  - Caches are smaller than main memory
  - Different cache levels are used with the quicker caches, like L1, have faster speeds but smaller sizes; L1 might be 256K, whereas the slower L2 might be 2M
- Data structures can have a significant impact on sparse matrix computation

Full Matrix versus Sparse Matrix Storage

- Full matrices are easily stored in arrays with just one variable needed to store each value since the value's row and column are implicitly available from its matrix position
- With sparse matrices two or three elements are needed to store each value
  - The zero values are not explicitly stored
  - The value itself, its row number and its column number
  - Storage can be reduced by storing all the elements in a particular row or column together
- Because large matrices are often quite sparse, the total storage is still substantially reduced

Sparse Matrix Usage Can Determine the Optimal Storage

- How a sparse matrix is used can determine the best storage scheme to use
  - Row versus column access; does the structure change
- Is the matrix essentially used only once? That is, its structure and values are assumed new each time used
- Is the matrix structure constant, with its values changed
  - This would be common in the N-R power flow, in which the structure doesn't change each iteration, but its values do
- Is the matrix structure and values constant, with just the **b** vector in **Ax=b** changing
  - Quite common in power system stability solutions

Numerical Precision

- Required numerical precision determines type of variables used to represent numbers
  - Specified as number of bytes, and whether signed or not
- For Integers
  - One byte is either 0 to 255 or -128 to 127
  - Two bytes is either smallint (-32,768 to 32,767) or word (0 to 65,536)
  - Four bytes is either Integer (-2,147,483,648 to 2,147,483,647) or Cardinal (0 to 4,294,967,295)
    - This is usually sufficient for power system row/column numbers
  - Eight bytes (Int64) if four bytes is not enough

Numerical Precision, cont.

- For floating point values using choice is between four bytes (single precision) or eight bytes (double precision); extended precision has ten bytes
  - Single precision allows for 6 to 7 significant digits
  - Double precision allows for 15 to 17 significant digits
  - Extended allows for about 18 significant digits
  - More bytes requires more storage
  - Computational impacts depend on the underlying device; on PCs there isn't much impact; GPUs can be 3 to 8 times slower for double precision
- For most power problems double precision is best

General Sparse Matrix Storage

- A general approach for storing a sparse matrix would be using three vectors, each dimensioned to number of elements
  - AA: Stores the values, usually in power system analysis as double precision values (8 bytes)
  - JR: Stores the row number; for power problems usually as an integer (4 bytes)
  - JC: Stores the column number, again as an integer
- If unsorted then both row and column are needed
- New elements could easily be added, but costly to delete
- Unordered approach doesn’t make for good computation since elements used next computationally aren’t necessarily nearby
- Usually ordered, either by row or column

Sparse Storage Example

- Assume

$$\mathbf{A} = \begin{bmatrix} 5 & 0 & 0 & -4 \\ 0 & 4 & 0 & -3 \\ 0 & 0 & 3 & -2 \\ -4 & -3 & -2 & 10 \end{bmatrix}$$

- Then

$$\mathbf{AA} = \begin{bmatrix} 5 & -4 & 4 & -3 & 3 & -2 & -4 & -3 & -2 & 10 \end{bmatrix}$$

$$\mathbf{JR} = \begin{bmatrix} 1 & 1 & 2 & 2 & 3 & 3 & 4 & 4 & 4 & 4 \end{bmatrix}$$

$$\mathbf{JC} = \begin{bmatrix} 1 & 4 & 2 & 4 & 3 & 4 & 1 & 2 & 3 & 4 \end{bmatrix}$$

Note, this example is a symmetric matrix, but the technique is general

Compressed Sparse Row Storage

- If elements are ordered (as was case for previous example) storage can be further reduced by noting we do not need to continually store each row number
- A common method for storing sparse matrices is known as the Compressed Sparse Row (CSR) format
  - Values are stored row by row
  - Has three vector arrays:
    - AA: Stores the values as before
    - JA: Stores the column index (done by JC in previous example)
    - IA: Stores the pointer to the index of the beginning of each row

CSR Format Example

- Assume as before

$$\mathbf{A} = \begin{bmatrix} 5 & 0 & 0 & -4 \\ 0 & 4 & 0 & -3 \\ 0 & 0 & 3 & -2 \\ -4 & -3 & -2 & 10 \end{bmatrix}$$

- Then

$$\mathbf{AA} = \begin{bmatrix} 5 & -4 & 4 & -3 & 3 & -2 & -4 & -3 & -2 & 10 \end{bmatrix}$$

$$\mathbf{JA} = \begin{bmatrix} 1 & 4 & 2 & 4 & 3 & 4 & 1 & 2 & 3 & 4 \end{bmatrix}$$

$$\mathbf{IA} = \begin{bmatrix} 1 & 3 & 5 & 7 \end{bmatrix}$$

CSR Comments

- The CSR format reduces the storage requirements by taking advantage of needing only one element per row
- The CSR format has good advantages for computation when using cache since (as we shall see) during matrix operations we are often sequentially going through the vectors
- An alternative approach is Compressed Sparse Column (CSC), which identical, except storing the values by column
- It is difficult to add values.
- We’ll mostly use the linked list approach here, which makes matrix manipulation simpler

Linked Lists: Classes and Objects

- In explaining the linked list approach it is helpful to use the concepts from object oriented programming (OOP) of classes and objects
  - Approach can also be used in non-OOP programming
- OOP can be thought of as a collection of objects interacting with each other
- Objects are instances of classes.
- Classes define the object fields and actions (methods)
- We'll define a class called sparse matrix element, with fields of value, column and next; each sparse matrix element is then an object of this class

Linked Lists

- A linked list is just a group of objects that represent a sequence
  - We'll used linked lists to represent a row or column of a sparse matrix
- Each linked list has a head pointer that points to the first object in the list
  - For our sparse matrices the head pointer will be a vector of the rows or columns

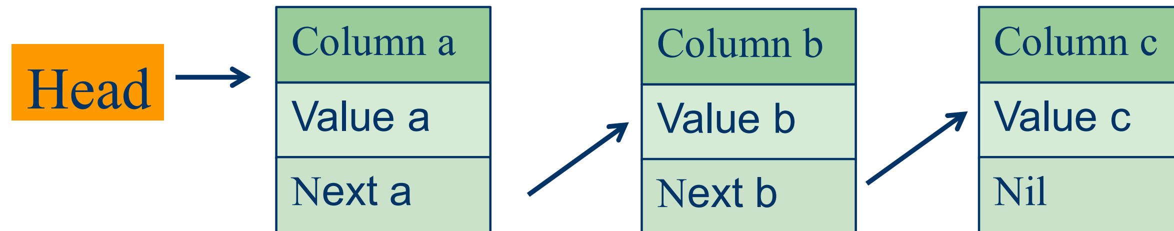

Sparse Matrix Storage with Linked Lists by Rows

- If we have an n by n matrix, setup a class called TSparseElement with fields column, value and next
- Setup an n-dimensional head pointer vector that points to the first element in each row
- Each nonzero corresponds to an object of class (type) TSparseElement
- We do not need to store the row number since it is given by the object's row
- For power system sparse matrices, which have nonzero diagonals, we also have a header pointer vector that points to the diagonal objects

Linked Lists, cont.

- Linked lists can be singly linked, which means they just go in one direction (to the next element), or doubly linked, pointing to both the previous and next elements
  - Mostly we'll just need singularly linked lists
- With linked lists it is quite easy to add new elements to the list. This can be done in sorted order just by going down the list until the desired point is reached, then changing the next pointer for the previous element to the new element, and for the new element to the next element (for a singly linked list)

Example Pascal Code for Writing a Sparse Matrix

```
Procedure TSparMat.SMWriteMatlab(Var ft : Text; variableName : String; digits,rod : Integer;
                              ignoreZero : Boolean; local_MinValue : Double);
Var j : Integer;
   p1 : TMatEle;
Begin
For j := 1 to n Do Begin
   p1 := Row(j).Head;
   While p1 <> nil Do Begin
     If (not IgnoreZero) or (abs(p1.value) > local_MinValue) Then Begin
       If variableName <> '' Then Writeln(ft,variableName+'(',(j),',',(p1.col),')=',p1.value:digits:rod,';')
       Else Writeln(ft,j:5,' ',p1.col:5,' ',p1.value:digits:rod);
     End;
     p1 := p1.next;
   End;
 End;
End;
```

Sparse Working Row

- Before showing a sparse LU factorization it is useful to introduce the concept of a working row full vector
- This is useful because sometimes we need direct access to a particular value in a row
- The working row approach is to define a vector of dimension n and set all the values to zero
- We can then load a sparse row into the vector, with computation equal to the number of elements in the row
- We can then unload the sparse row from the vector by storing the new values in the linked list, and resetting the vector values we changed to zero

Loading the Sparse Working Row

```
Procedure TSparMat.LoadSWRbyCol(rowJ : Integer; var SWR : PDVectorList);
Var p1 : TMatEle;
Begin
  p1 := rowHead[rowJ];
  While p1 <> nil Do Begin
    SWR[p1.col] := p1.value;
    p1 := p1.next;
  End;
End;
```

Unloading the Sparse Working Row

```
Procedure TSParMat.UnLoadSWRbyCol(rowJ : Integer; var SWR :
PDVectorList);
Var p1 : TMatEle;
Begin
  p1 := rowHead[rowJ];
  While p1 <> nil Do Begin
    p1.value := SWR[p1.col];
    SWR[p1.col] := 0;
    p1 := p1.next;
  End;
End;
```

Note, there is no need to explicitly zero out all the elements each iteration since 1) most are still zero and 2) doing so would make it $O(n^2)$. The above code efficiently zeros out just the values that have changed.

Doing an LU Factorization of a Sparse Matrix with Linked Lists

- Now we can show how to do an LU factorization of a sparse matrix stored using linked lists
- We will assume the head pointers are in the vector RowHead, and the diagonals in RowDiag
- Recall this was the approach for the full matrix

```
For i := 2 to n Do Begin  // This is the row being processed
 For j := 1 to i-1 Do Begin  // Rows subtracted from row i
  A[i,j] = A[i,j]/A[j,j]  // This is the scaling
  For k := j+1 to n Do Begin  // Go through each column in i
   A[i,k] = A[i,k] - A[i,j]*A[j,k]
  End;
 End;
End;
```

Sparse Factorization

- Note, if you know about fills, we will get to that shortly; if you don't know don't worry about it yet
- We'll just be dealing with structurally symmetric matrices (incidence-symmetric)
- We'll assume the row linked lists are ordered by column; we'll show how this can be done quickly later
- We will again sequentially going through the rows, starting with row 2, going to row n

For i := 2 to n Do Begin // This is the row being processed

Sparse Factorization, cont.

- The next step is to go down row i, up to but not including the diagonal element
- We’ll be modifying the elements in row i, so we need to load them into the working row vector
- Key sparsity insight is in doing the below code we only need to consider the non-zeros in A[i,j]; for a full matrix the code is

```
For j := 1 to i-1 Do Begin  // Rows subtracted from row
    A[i,j] = A[i,j]/A[j,j]  // This is the scaling
    For k := j+1 to n Do Begin  // Go through each column in i
            A[i,k] = A[i,k] - A[i,j]*A[j,k]
    End;
```

Sparse Factorization, cont.

```
For i := 1 to n Do Begin   // Start at 1, but nothing to do in first row
  LoadSWRbyCol(i,SWR);   // Load Sparse Working Row }
  p2 := rowHead[i]
  While p2 <> rowDiag[i] Do Begin    // This is doing the j loop
     p1 := rowDiag[p2.col];
     SWR[p2.col] := SWR[p2.col] / p1.value;
     p1 := p1.next;
     While p1 <> nil Do Begin   // Go to the end of the row
       SWR[p1.col] := SWR[p1.col] - SWR[p2.col] *p1.value;
       p1 := p1.next;
     End;
     p2 := p2.next;
   End;
   UnloadSWRByCol(i,SWR);
 End;
```

Sparse Factorization Example

- Believe it or not, that is all there is to it! The factorization code itself is quite simple.
- However, there are a few issues we'll get to in a second. But first an example

$$\mathbf{A} = \begin{bmatrix} 5 & 0 & 0 & -4 \\ 0 & 4 & 0 & -3 \\ 0 & 0 & 3 & -2 \\ -4 & -3 & -2 & 10 \end{bmatrix}$$

- Notice with this example there is nothing to do with rows 1, 2 and 3 since there is nothing before the diag (p2 will be equal to the diag for the first three rows)

Sparse Factorization Example, Cont.

- Doing factorization with i=4
  - Row 4 is full so initially p2= A[4,1] // column 1
  - SWR = [-4 -3 -2 10]
  - p1= A[1,1]
  - SWR[1] = -4/A[1,1] = -4/5 = -0.8
  - p1 goes to A[1,4] That is, the next element in row 1
  - SWR[4] = 10 – SWR[p2.col]*p1.value = 10 – (-0.8)*-4=6.8
  - p1 = nil; go to next col
  - p2 =A[4,2] // column 2
  - P1 = A[2,2]
  - SWR[2] = -3/A[2,2]= -3/4 = -0.75

Sparse Factorization Example, Cont.

- p1 goes to A[2,4] // p2=A[4,2]   The next element in row 2
- SWR[4] = 6.8 – SWR[p2.col]*p1.value = 6.8 – (-0.75)*-3=4.55
- p1 = nil; go to next col
- p2 =A[4,3] // column 3
- p1 = A[3,3]
- SWR[3] = -/A[2,2]= -2/3 = -0.667
- p1 goes to A[3,4] // p2 = A[4,3]   The next element in row 3
- SWR[4] = 4.55 – SWR[p2.col]*p1.value
  = 4.55 – (-0.667)*-2=3.2167
- Unload the SWR = [-0.8 -0.75 -0.667 3.2167]
- p2 = A[4,4] = diag so done

Sparse Factorization Examples, Cont.

$$\mathbf{A}_{\textbf{Factored}} = \begin{bmatrix} 5 & 0 & 0 & -4 \\ 0 & 4 & 0 & -3 \\ 0 & 0 & 3 & -2 \\ -0.8 & -0.75 & -0.6667 & 3.2167 \end{bmatrix}$$

- For a second example, again consider the same system, except with the nodes renumbered

$$\mathbf{B} = \begin{bmatrix} 10 & -4 & -3 & -2 \\ -4 & 5 & 0 & 0 \\ -3 & 0 & 4 & 0 \\ -2 & 0 & 0 & 3 \end{bmatrix}$$

Sparse Factorization Examples, Cont.

- With i=2, load SWR = [-4 5 0 0]
  - p2 = B[2,1]
  - p1 = B[1,1]
  - SWR[1]=-4/p1.value=-4/10 = -0.4
  - p1 = B[1,2]
  - SWR[2]=5 – (-0.4)*(-4) = 1.6
  - p1 = B[1,3]
  - SWR[3]= 0 – (-0.4)*(-3) = -1.2
  - p1 = B[1,4]
  - SWR[4]=0 – (-0.4)*(-2) = -0.8
  - p2=p2.next=diag so done
  - UnloadSWR and **we have a problem!**

$$\mathbf{B} = \begin{bmatrix} 10 & -4 & -3 & -2 \\ -4 & 5 & 0 & 0 \\ -3 & 0 & 4 & 0 \\ -2 & 0 & 0 & 3 \end{bmatrix}$$

There are no elements in row 2 for columns 3 and 4!

Fills

- When doing a factorization of a sparse matrix some values that were originally zero can become nonzero during the factorization process
- These new values are called "fills"
  (some call them fill-ins)
- For a structurally symmetric matrix the fill occurs for both the element and its transpose value (i.e., $\mathbf{A}_{ij}$ and $\mathbf{A}_{ji}$)
- How many fills are required depends on how the matrix is ordered
  - For a power system case this depends on the bus ordering

Fills

- There are two key issues associated with fills
  - Adding the fills
  - Ordering the matrix elements (buses in our case) to reduce the number of fills
- The amount of computation required to factor a sparse matrix depends upon the number of nonzeros in the original matrix, and the number of fills added
- How the matrix is ordered can have a dramatic impact on the number of fills, and hence the required computation
- Usually a matrix cannot be ordered to totally eliminate fills

Fill Examples

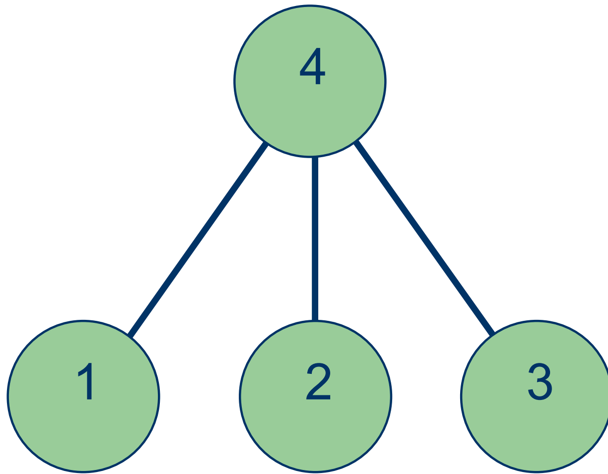

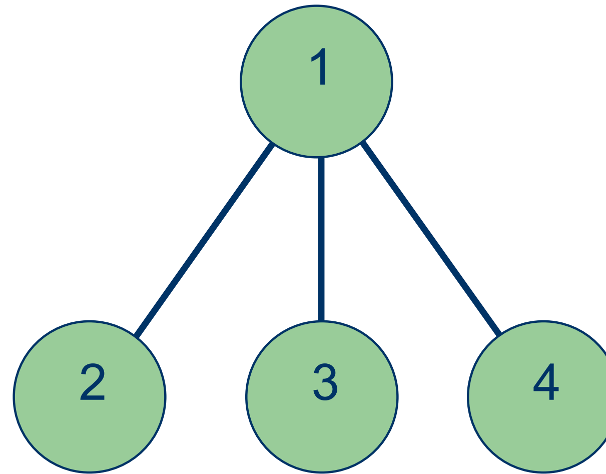

$$\mathbf{A} = \begin{bmatrix} 5 & 0 & 0 & -4 \\ 0 & 4 & 0 & -3 \\ 0 & 0 & 3 & -2 \\ -4 & -3 & -2 & 10 \end{bmatrix}$$

No Fills Required

$$\mathbf{B} = \begin{bmatrix} 10 & -4 & -3 & -2 \\ -4 & 5 & 0 & 0 \\ -3 & 0 & 4 & 0 \\ -2 & 0 & 0 & 3 \end{bmatrix}$$

Fills Required (matrix becomes full)

Example: 7 by 7 Matrix

- Consider the 7 x 7 matrix **A** with the zero-nonzero pattern shown in (*a*): of the 49 possible elements there are only 31 that are nonzero
- If elimination proceeds with the given ordering, all but two of the 18 originally zero entries, will fill in, as seen in (*b*)

Example: 7 by 7 Matrix Structure

r \ c	1	2	3	4	5	6	7
1	X	X	X	X	X	X	
2	X	X	X			X	X
3	X	X	X			X	X
4	X			X	X		
5	X			X	X	X	
6	X	X	X		X	X	
7		X	X				X

The original zero-nonzero structure

r \ c	1	2	3	4	5	6	7
1	X	X	X	X	X	X	
2	X	X	X	F	F	X	X
3	X	X	X	F	F	X	X
4	X	F	F	X	X	F	F
5	X	F	F	X	X	X	F
6	X	X	X	F	X	X	F
7		X	X	F	F	F	X

The post- elimination zero nonzero pattern

Example: 7 by 7 Matrix Reordering

- We next reorder the rows and the columns of **A** so as to result in the pattern shown in (c)
- For this reordering, we obtain no fills, as shown in the table of factors given in (d)
- In this way, we preserve the original sparsity of **A**

Example: 7 by 7 Matrix Reordered Structure

r \ c	4	5	1	6	7	3	2
4	X	X	X				
5	X	X	X	X			
1	X	X	X	X		X	X
6		X	X	X		X	X
7					X	X	X
3			X	X	X	X	X
2			X	X	X	X	X

The reordered system

r \ c	4	5	1	6	7	3	2
4	X	X	X				
5	X	X	X	X			
1	X	X	X	X		X	X
6		X	X	X		X	X
7					X	X	X
3			X	X	X	X	X
2			X	X	X	X	X

The post- elimination reordered system

Graph (Grid) Insights

- For electric grids it can be helpful to relate the sparse matrix back to an associated electric grid
- Assume an n by n matrix A
- All the diagonals are non-zeros
- If there is a connection between two buses (nodes), say at position j and k, then the associated matrix entries, $\mathbf{A}_{jk}$ and $\mathbf{A}_{kj}$ are nonzero